

Region Based Approximation for High Dimensional Bayesian Network Models

Peng Lin, Martin Neil, Norman Fenton

Abstract—Performing efficient inference on Bayesian Networks (BNs), with large numbers of densely connected variables is challenging. With exact inference methods, such as the Junction Tree algorithm, clustering complexity can grow exponentially with the number of nodes and so computation becomes intractable. This paper presents a general purpose approximate inference algorithm called Triplet Region Construction (TRC) that reduces the clustering complexity for factorized models from worst case exponential to polynomial. We employ graph factorization to reduce connection complexity and produce clusters of limited size. Unlike MCMC algorithms TRC is guaranteed to converge and we present experiments that show that TRC achieves accurate results when compared with exact solutions.

Index Terms—Belief propagation, High dimensional, Bayesian Networks, Graph factorization, discrete energy optimization

1 MOTIVATION AND CONTRIBUTION

PERFORMING efficient inference on Bayesian network (BN) models with a large number of variables that are also densely connected (high order dependent), is a major computational challenge. With exact methods, such as the Junction Tree (JT) algorithm [1], [2], the complexity depends on the size of maximal cluster of the triangulated graph [2], [3], and the maximal cluster size can grow exponentially with the number of nodes.

An important way to reduce this complexity is to apply factorization algorithms, such as binary factorization [4], to convert the dense model into a factorized model resulting in reduced connection complexity. However, the resulting factorized BN's tree-width¹ remains high, and cluster size in the factorized BN remains exponential as clustering will result in a product of all factors involved in the maximum cluster.

As exact inference on high tree-width models is generally intractable, there are several techniques that use the bounded tree-width of JTs, such as [5], [6]. These so called thin-JTs ensure tractability by using an upper bound of the tree-width. However, the overall performance on accuracy is not guaranteed, except where child nodes are deterministically related to their parents or where the modeller can make context-specific independence assumptions [7].

When exact methods cannot be performed efficiently on high tree-width models, sampling based methods, such as Markov Chain Monte Carlo (MCMC) [1], [8], are used, but MCMC solutions usually have to be tailored to the problem and convergence is not guaranteed.

Motivated by the success of using region based approximation for spin class/grid models [9], we employ region based approximation to reduce the clustering complexity for high tree-width factorized models [4], [10]. We present a general inference algorithm called Triplet Region Construction (TRC) based on region belief propagation and show how it can perform robust inference on BNs. In doing so a

number of existing well known challenges (region choice, convergence and accuracy) encountered when using region based approximation are addressed. Most significantly our methods provide an algorithm where the clustering and efficiency complexity for factorized models is reduced from worst case exponential to polynomial.

The paper is structured as follows:

In section 2, we introduce necessary BN notation and background and explain why it is sufficient for our proposed algorithm to focus on complete BN models via a uniquely defined binary factorized model which has lower order dependences.

In section 3, we introduce region based approximation methods and explain why previous region based algorithms cannot deal with the factorized models defined in Section 2. In section 4, we show how region based approximation can be used to develop our proposed TRC algorithm and reduce the clustering complexity.

In section 5, we present experiments involving sparse BN and high tree-width factorized BNs to show the accuracy and robustness of TRC. We also contrast these results with those obtained using MCMC and discuss the accuracy achieved.

Section 6 concludes the paper and discusses extensions of the TRC algorithm.

The paper provides four key novel contributions:

- 1) Whereas region based belief propagation [9] is typically applied to undirected graphical models here we present its use, in a systematic way, for directed models for the first time (sections 3 and 4).
- 2) Construction of a region graph for general models is difficult because the choice of regions and interactions is left up to the model designer, with varied results. We present a region identification algorithm, called Outer Region Identification (ORI), that incorporates all local (considered conditional independence) factor correlations as an effective way of identifying the largest regions (sections 4.1 and 4.2).

1. Tree-width is one less than the minimum possible value of the maximum cluster membership size over all possible triangulations.

ORI can be used separately to provide the region specification for many other region graphs. We then use redundant regions resulting from ORI, to adjust our region graph to satisfy the perfect correlation property² [9], [11] and maxent-normal property³ [9] that are necessary conditions for computational accuracy.

- 3) Previous region based algorithms suffer from unavoidable numerical instability problems when performing inference on high tree-width factorized models. We propose a Region Graph Binary Factorization (RGBF) algorithm to decompose the region graph into an equivalent, but more numerically stable, alternative. We show that RGBF improves the robustness of region based belief propagation algorithms (section 4.4). RGBF is also a separate algorithm that can be used for any region graph.
- 4) Finally (and most importantly) we describe the TRC algorithm (section 4.5) in terms of the above sub-algorithms. TRC is guaranteed to converge, solves the complexity challenge encountered in clustering high tree-width factorized models and achieves accurate results when we compare the marginal distributions of individual variables with those produced by JT and with MCMC (section 5).

2 COMPLETE BN AND ITS DECOMPOSITION

In this section we provide a brief overview of BNs and their notation and then discuss why it is sufficient to solve the BN inference problem using a complete BN graph, and how a BN's connection complexity can be reduced using binary factorization.

A BN is a directed acyclic graph (DAG), with nodes X_1, X_2, \dots, X_n representing random variables (which can be discrete or continuous⁴), together with a conditional probability distribution (CPD) for each node which is conditional on its parent nodes if there are any (for discrete variables we refer to node probability tables (NPTs)). The absence of arcs between nodes encodes the Conditional Independence (CI) [1] assumptions between variables. The BN represents the joint distribution, p of the random variables X_1, X_2, \dots, X_n as the product of its CPDs. In the absence of CI assumptions, we can use the chain rule to factorize the joint distribution, as shown in equation (1).

$$p(X_1, X_2, \dots, X_n) = \prod_{i=1}^n p(X_i | X_1, \dots, X_{i-1}) \quad (1)$$

With CI assumptions this simplifies to

$$P(X_1, \dots, X_n) = \prod P(X_i | pa(X_i))$$

where $pa\{X_i\}$ represents the parents of node X_i . This simplification, along with the associated graphical

representation, is one of the attractions of using BNs.

However, in the worst case there are no CI assumptions in the BN. In this case the BN graph is a *complete* DAG with n nodes, i.e. every pair of nodes is connected by a directed edge. Performing inference on such a complete BN graph represents the worst case complexity for exact algorithms and is usually intractable. Since it is therefore assumed to be impossible to find efficient exact algorithms for arbitrary BN models, the challenge is to find good approximate algorithms.

It is crucial to note that any BN model can be regarded as a complete graph with some edges removed, where the remaining edges encode the CI assumptions. Conversely, any non-complete BN graph model (referred to as a sparse graph) can be converted to a complete graph model by adding appropriate edges. So, theoretically, any BN model can be represented by a complete graph. Hence, any algorithm that performs efficient inference for complete BN models will also be efficient for arbitrary BN models. It is therefore sufficient to find an approximate inference algorithm that is efficient for complete BN models and hence, this is the focus for the rest of the paper.

In what follows we also make use of a well known result of graph theory (which can be proved by induction on the number of nodes) which asserts that any complete DAG of n nodes has a unique Hamiltonian path, and is hence uniquely defined up to a permutation of the n nodes. Specifically,

Theorem 1: In any complete DAG of n nodes there is exactly one node with indegree $n-1$, exactly one node with indegree $n-2, \dots$, exactly one node with indegree one, and exactly one node with indegree zero.

Theorem 1 ensures the uniqueness of the chain rule factorization (Equation 1) for a complete DAG of n nodes subject to the order in which, for each $i = 1, \dots, n$ node X_i is the (unique) node with indegree $i-1$. In what follows we will assume this ordering of the nodes in the complete graph.

In addition to assuming a complete BN model we also need to transform the complete BN graph model into a version that is binary factorized (and equivalent in the sense defined below) by introducing additional nodes in such a way that each node has at most two parents. We call the process of producing a Binary Factorized BN (a BFG), the BF-process. A BFG avoids the computational complexity problem of exponential size CPDs (although, at this stage, the cluster size is not reduced). Instead each CPD has at most three members, i.e. is at worst a triplet. Furthermore, we will also benefit from this complexity reduction when applying the region based approximation discussed in section 4.

Proposition 1: A BN G can be transformed into a binary factorized BN G' (i.e. each node has at most two parents) whose nodes are a superset of G and which is 'equivalent' to G in the sense that, for each node X in G , the CPD of X in G' is equivalent after factorization to the CPD of X in G .

Proof. In what follows we assume the unique ordering of the

2. This requires the sum of all regions' counting numbers to be one, i.e. $\sum_R c_R = 1$ which ensures that the region-based entropy is correct if all variables in p are perfectly correlated.

3. A constrained region-based free energy approximation is maxent-normal if it is valid and the corresponding constrained region-based entropy H_G achieves its maximum when all the beliefs are uniform.

4. All continuous variables can be assumed to be discretized statically or dynamically [3], [10], [12], [13], [14].

complete graph G from Theorem 1, and apply the structural factorization of G as described in [4], by introducing a set of intermediate variables E_t ($E_t \in G', E_t \notin G$) that are not in the original BN. For example, in the case of the 5-dimensional complete graph, the structure of the binary factorized version is as shown in Figure 1. While this BF algorithm is guaranteed to produce a uniquely structured BFG G' or each complete BN graph G , we have to show how to define the CPDs in G' so that for each node X in G , the CPD of X in G' is equivalent after factorization to the CPD of X in G . There are three types of nodes whose CPDs we have to consider: 1) Continuous nodes with continuous parents only; 2) Discrete nodes with discrete parents only; 3) Mixture nodes (Continuous nodes with at least one discrete parent or discrete nodes with at least one continuous parent):

- 1) **Continuous nodes with continuous parents:** The result for this case was proved in [4]. In summary, we assume a continuous CPD $P(Z|pa\{Z\})$ is always expressed as an arithmetical expression over Z and $pa\{Z\}$, and this expression can always be parsed incrementally by smaller expressions which involve only two variables. By introducing the intermediate nodes, this naturally results in a BF process for the continuous case. The BF process in [4] was motivated to ensure the same equivalence property as is required here. Figure 1, provides an example applied to the 5-dimensional complete graph in the simple case where all nodes are continuous linear functions of their parents. In the resulting BFG G' (with the exception of the root nodes X_1 and X_2), each node has exactly two parents.
- 2) **Discrete nodes with discrete parents:** In general a discrete node D with three discrete parents A , B and C can be transformed into an equivalent binary factorised form by introducing an intermediate node E (with parents A and B) that has $n \times m$ states e_{ij} ($i = 1, \dots, n$ and $j = 1, \dots, m$) where A has n states a_1, \dots, a_n and B has m states b_1, \dots, b_m . The NPT for E is defined as:

$$P(E = e_{ij}|a_k, b_l) = \begin{cases} 1 & \text{if } k = i \text{ and } l = j \\ 0 & \text{otherwise} \end{cases}$$

The NPT for node D in G' (with parents C and E) is defined as:

$$P(D|e_{ij}, c_k) = P_G(D|a_i, b_j, c_k)$$

Figure 2 shows the full solution for the 5-dimensional complete graph. The method is applied iteratively when there are more than three discrete parents.

- 3) **Mixture nodes:** Suppose Z is a mixture node. We consider the two cases:
 - a) Z is a continuous node with at least one discrete parent. If Z has more than one discrete parent then we can apply the BF process described above for the discrete node parents, to ensure an equivalent factorization of those nodes such that Z has just one discrete parent Y . So we can assume Z has exactly one discrete parent Y and that the CPD

for Z is $P(Z|pa\{Z\}) = \sum P(Y = y_i) \cdot f(X_i)$ ($Y, X_i \in pa\{Z\}$), where X_i are continuous variables. Then Z can be binary factorized by incrementally combining the X_i densities, $i > 2$. The CPD for the intermediate variable E_k can be defined as:

$$P(E_k|pa\{E_k\}) = \frac{P(Y = y_i) \cdot X_i + P(Y = y_j) \cdot X_j}{P(Y = y_i) + P(Y = y_j)}$$

($X_i, X_j \in pa\{E_k\}$). The CPD of Z can be recovered by incrementally combining E_k with another parent X_q ($q \neq i \neq j$) by using the same formula as defined for $P(E_k|pa\{E_k\})$.

b) Z is a discrete node with at least one continuous parent. In this case we can apply the BF process for continuous nodes to factorize on its parent nodes and guarantee Z has only two parents. An example of the mixture node case is shown in Figure 3. Specifically, X_1 is a discrete node and the other nodes are continuous; for example, X_5 is a mixture node with one discrete parent and three continuous parents. \square

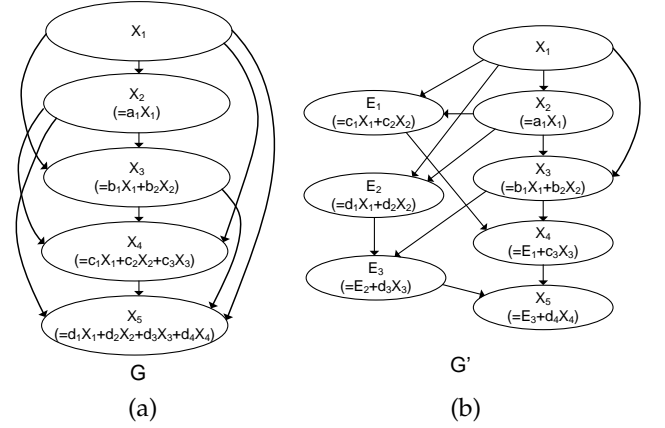


Fig. 1. (a) 5 dimensional complete graph G with all continuous variables; (b) resulting BFG G'

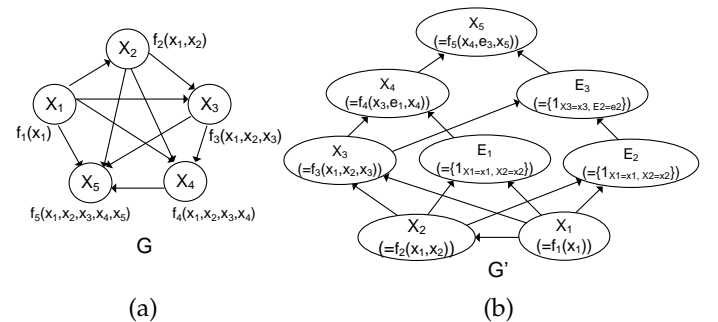


Fig. 2. (a) 5 dimensional complete graph G with all discrete variables, nodes X_1, \dots, X_5 are associated to CPDs f_1, \dots, f_5 ; (b) resulting BFG G'

Applying the BF process to an n node complete BN graph model results in a BFG model with κ_n nodes where

$$\kappa_n = n + (n-2)(n-3)/2 = (n^2 - 3n + 6)/2$$

We will use κ_n to denote the number of variables in a BFG throughout the rest of the paper.

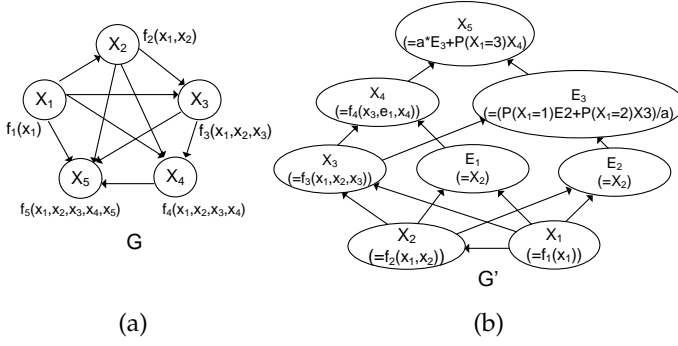


Fig. 3. (a) 5 dimensional complete BN G with nodes X_1, \dots, X_5 and associated CPDs f_1, \dots, f_5 ; (b) BF process for (a) if (a) is mixture with X_1 as the discrete variable, and others the continuous variables. The CPD for X_5 is defined as: $f_5 = P(X_1 = 1) \cdot X_2 + P(X_1 = 2) \cdot X_3 + P(X_1 = 3) \cdot X_4$, with constant $\alpha = P(X_1 = 1) + P(X_1 = 2)$. The function f_4 is not factorized but only has changed variables as only two continuous parent variables are involved.

3 REGION BASED BELIEF PROPAGATION

As already noted, while binary factorization results in reduced factor size complexity and is a necessary step in the TRC algorithm that we propose, it does not solve the problem of cluster size complexity. Hence it does not avoid the complexity problems associated with exact algorithms, such as JT. When an n -dimensional complete BN is converted to a non-complete BN, the tree-width remains $n - 1$ in the factorized BN. This is also evident in a $D \times D$ grid where the model is sparse but the tree-width is $O(D)$ [15]. In these circumstances the triangulation graph produces clusters that are still (in the worst case) exponential.

In section 3.1 we will give an overview of region based approximation [7], [9], [16], [17] which partly addresses the problem of cluster size complexity in our TRC algorithm, and section 3.2 discusses how it presents residual difficulties for high tree-width BN model problems (which we subsequently resolve in Section 4).

3.1 Review of GBP and CCCP Algorithms

Before introducing region based belief propagation, we define region and region graph.

Definition 1. A region r of a factor graph is a set of variable nodes V_r and factor nodes A_r , such that if a factor node a is in A_r , all variable nodes neighbouring a are in A_r . [9]

Definition 2. A region graph \mathcal{G} is a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{L})$ in which each vertex $v \in \mathcal{V}$ corresponding to a region r is labelled (denote the label of vertex v by $l(v) \in \mathcal{L}$) with a subset of nodes in a factor graph. We say v_p is a parent of v_c if $v_p \rightarrow v_c$ is a directed edge $e \in \mathcal{E}$. [9]

Region based belief propagation [9] is popularly used in specific undirected graphic models, such as finite dimensional spin class models [9], [18], [29] encoded by a factor graph⁵ These models may be intractable for exact methods but can be approximated using region

based belief propagation. Here the clustering complexity reduction is achieved by constructing variational region based complexities and involves the approximation of a free energy function term and its function space.

Performing exact inference is equivalent to solving an optimization problem over the exact energy function $F[p, Q]$, where p is a distribution over χ (the set of all possible assignments of values to all the network's random variables) and Q is the space of all marginals. This is NP-hard [7], [18]. When $F[p, Q]$ cannot be tractably optimized a factored energy function can be defined in terms of entropies over all regions in a region graph \mathcal{G} , as an approximation of $F[p, Q]$. However, even approximation of $F[p, Q]$ over the marginal polytope $\text{marg}[\mathcal{G}]$ ($\text{marg}[\mathcal{G}] = \{Q_p\}$) is also an NP-hard problem [7]. Instead, we perform optimization over the locally consistent polytope $\text{local}[\mathcal{G}]$, which is a set of pseudo-marginal distributions over the variables in each region. This local consistency achieves a polynomial computation complexity provided that all regions are calibrated and neighbouring regions are locally consistent with each other [7].

Yedidia et al. [9] demonstrated that the convergence of these self-consistent constrained region belief equations corresponds to the minimal points of the Kikuchi variation free energy. Minimization of the Kikuchi cluster free energy is equivalent to the problem of constructing a fixed point for the region belief equations and this can be achieved using iterative message-passing over a region graph. In this section, we review two of the region based message-passing algorithms [9], [16], [19], [20] that we use as the starting point for our approach. The first is Generalized Belief Propagation (GBP) [9], which is a generalization of a class of belief propagation based algorithms, involving Loopy Belief Propagation [21], Survey Propagation [22], and others. However, GBP does not guarantee convergence⁶, and although the second message passing algorithm, called the Concave Convex Procedure (CCCP) [20], [24], does guarantee convergence it can, unfortunately, be numerically unstable for large models. Indeed, performing inference on the region graph for our high tree-width BFG models using both algorithms is numerically unstable and/or does not converge. Other BP related algorithms, such as expectation propagation [16], non-parametric belief propagation [25], particle belief propagation [26], are not discussed here as they characterize messages in a continuous domain only with additional approximations and also do not guarantee convergence.

For convenience to use GBP we can apply the Cluster Variation Method (CVM) [9] to produce a valid region graph. Firstly the outer regions⁷ are identified for the first level of a region graph and next CVM generates the regions for subsequent levels using intersections of the regions declared at the previous level. The Kikuchi region based free energy function $F_{\mathcal{G}}$ is defined in Equation 2. Our task is to minimize $F_{\mathcal{G}}$ under a set of self-consistent constraints

6. It is possible to use the so-called dumping technique [9], [23] to help converge but it is still not guaranteed.

7. An outer region is defined to be those regions with no parents, i.e. having no incoming region edges.

5. A bipartite graph representing the factorization of a function, with factor node containing all factors.

imposed by the CVM region graph.

$$F_G = \sum_{r \in R} c_r \left\{ \sum_{x_r} b_r(x_r) E_r(x_r) + \sum_{x_r} b_r(x_r) \log b_r(x_r) \right\} + L_G \quad (2)$$

where R is the set of all regions in \mathcal{G} , $c_r = 1 - \sum_{r' \in \text{Ancestor}(r)} c_{r'}$, and corresponds to counting number of each region, given $c_{r'}$ is number of degrees of freedom for the region. The term $E_r(\mathbf{x})$ represents the energy associated with region r . The region belief term b_r is an estimated distribution of the true distribution over region r [9].

To perform minimization of Equation 2 we need to consider the Lagrangian term L_G (shown in equation 3), which incorporates two kinds of constraints: the normalization constraint for each region, $\sum_{x_r} b_r(x_r) = 1$ and the running intersection constraints between parent and child region beliefs, $\sum_{x \in r \setminus c} b_r(x_r) = b_c(x_c)$ ($c \in \text{child}(r)$).

$$L_G = \sum_{r \in R} \sum_{c \in \text{child}(r)} \sum_{x_c} \lambda_{r,c}(x_c) \left\{ \sum_{x \in r \setminus c} b_r(x_r) - b_c(x_c) \right\} + \sum_{r \in R} \gamma_r \left(\sum_{x_r} b_r(x_r) - 1 \right) \quad (3)$$

Solving for Lagrangian multipliers λ and γ , in (3), corresponds to an iterative message passing algorithm. Both GBP and CCCP updated equations involve solving Lagrangian multipliers. The difference between CCCP and GBP is that CCCP splits the free energy function F_G into a concave and a convex part, $F_G = F_{G,\text{cave}} + F_{G,\text{vex}}$ and GBP does not. Minimization of the free energy function is then an iterative procedure guaranteed to minimize the convex part and maximize the concave part using tangent matching [20]. In general, GBP is more efficient than CCCP since CCCP uses an outer-inner double loop and each inner loop involves recursively updating the Lagrangian multipliers.

3.2 Region Based Approximation Difficulties for Directed Models

Region graphs generated for high tree-width directed models often involve multiple connections between regions located at different levels. Because one parent can have a large number of children the same variable will appear in many different regions across these levels. This gives rise to a large counting number and multiple cycles associated with a single region, leading to under/overflows during the multiplication of multiple messages. Such numerical instability might be encountered during message updating in both GBP and CCCP. The message updating process is shown in equations 4, 5 and 6 for GBP.

$$b_r(x_r) = \tilde{f}_r(x_r) \prod_{c \in \text{child}(r)} n_{c \rightarrow r}(x_c) \prod_{p \in \text{parent}(r)} m_{p \rightarrow r}(x_p) \quad (4)$$

$$m_{p \rightarrow r}(x_r) = (n_{r \rightarrow p}^0(x_r))^{\beta_r - 1} (m_{p \rightarrow r}^0(x_r))^{\beta_r} \quad (5)$$

$$m_{r \rightarrow c}^0(x_c) = \sum_{x_r \setminus x_c} \tilde{f}_r(x_r) \prod_{p \in \text{parent}(r)} m_{p \rightarrow r}(x_r) \prod_{c' \in \text{child}(r) \setminus c} n_{c' \rightarrow r}(x_{c'}) \quad (6)$$

where $\tilde{f}_r(\mathbf{x}_r) \equiv (\prod_{a \in A_r} f_a(\mathbf{x}_a))^{c_r}$, $\beta_r = 1/(2 - (1 - c_r)/p_r)$ (p_r is the number of parents of r) and $n_{c \rightarrow r}$ is similarly updated as $m_{p \rightarrow r}$ by pseudo-messages ($m_{r \rightarrow c}^0(x_c)$).

As all messages are exponential family distributions, a large absolute value of counting number $|c_r|$ will result in either large or small values of f_r , which in turn influences $m_{p \rightarrow r}$ and also b_r . The over/under flow problem is unavoidable when the number of dimensions is large. Likewise, a large counting number is caused by multiple paths from parents to the same child, resulting in many cycles in the region graph. Too many cycles in a region graph will make message scheduling difficult and inhibit convergence. The region graphs constructed for high tree-width BFG models inevitably encounter these large counting numbers and multiple cycles problems, and as a result both GBP and CCCP are indeed numerically unstable. In summary, there are three major difficulties preventing the use of region based belief propagation to approximate high tree-width BFG models:

- 1) The construction of region graphs has, to date, been ad-hoc and problem specific, thus making generalisation and accuracy difficult. We address this in sections 4.1 and 4.2.
- 2) It is not clear on how much interaction strength⁸ [27] among outer regions is needed. This problem is also addressed in sections 4.1 and 4.2.
- 3) By using CVM to generate regions for high tree-width BFG models the same variables can appear in multiple regions in the region graph, leading to multiple cycles associated with the smallest regions and numerical instability. This problem is addressed in section 4.4.

4 TRIPLET REGION CONSTRUCTION

This section summarises the TRC algorithm, designed to address the difficulties that arise when using region based approximation for high tree-width directed models. TRC is composed of three sub algorithms. In Section 4.1, we provide a formulation for converting a BN to a parametric Markov Network and we propose the Outer Region Identification Algorithm (ORI) to identify outer regions as a first step to construct a valid region graph. In Section 4.2, we illustrate how redundant outer regions are identified and can be rejected, so we can produce a region graph which will satisfy desired properties for accuracy. Section 4.3 proves that the region graph we proposed satisfies the desired properties for accuracy. To avoid numerical instability, in Section 4.4 we propose the Region Graph Binary Factorization (RGBF) algorithm to ensure each region has exactly two parent regions. Section 4.5 describes the TRC algorithm as a combination of ORI, RGBF and CCCP.

8. This is the number of interactions between outer regions.

4.1 Outer Region Identification Algorithm

Since belief propagation algorithms are typically designed for undirected models, such as Markov Networks⁹ (MNs), for directed models we need to convert the CPDs into factors.

To convert a BN to an undirected parameterization we first identify factors ϕ such that $\phi_{\{X_i\} \cup pa\{X_i\}}(X_i, pa\{X_i\}) = P(X_i|pa\{X_i\})$.

We also need to construct the *moral graph* of a BN G , denoted $M[G]$. This is an undirected graph that contains an edge (X_i, X_j) if there is an edge between X_i and X_j in G , or if X_i and X_j are parents of the same child node¹⁰. The added edge between the parents that share the same child node is called a *moral edge*.

Unfortunately, connecting the parent nodes X_i and X_j via a moral edge assumes that $X_i \not\perp X_j$ and we may lose CI information contained in the original BN (i.e. $I(M[G]) \subseteq I(G)$). To resolve this, the CI information will be incorporated during the region graph construction.

Constructing a good region graph is an open research question, because minimizing the free energy function is a necessary but not sufficient condition. For instance, Welling et al. [11], [27], [28] discuss ways to produce Structure Region Graphs based on graphical topology and they offer guidance based on structural information criteria. They define outer regions based on basic cycles in the MN graph. Interaction is optimal when the cycles can be ordered so that each cycle has some edge that does not appear in any cycle preceding it in the ordering. However, optimization from graphical topological information is difficult in the absence of some general rules.

Therefore, the choice of outer regions is a vital first step and if this is incorrect or suboptimal convergence and stability cannot be guaranteed.

We call our approach to identifying the outer regions **Outer Region Identification (ORI)**, which identifies outer regions in terms of maximizing interaction strength as well as reducing redundancy.

As p' is already represented by the product of only triplet factors derived by the BF algorithm, any factors larger than triplets can be decomposed into triplet factors. Due to the Markov property this is an exact procedure for the decomposition. As a result, we can include all triplet factors as outer regions at the first level of a CVM region graph, which will produce a valid region graph [9]. The remaining problem is to determine the interaction strength among these outer regions. For this we need to introduce the *Maximally Exhaustive* property.

Definition 3. A region graph satisfies the *Maximally Exhaustive* property if any maximum subset of the outer region that contains a factor converted from the original BN is included in at least one second level region.

Theorem 2: In a valid CVM region graph for BFG models with all triplet factors as outer regions interaction strength

is sufficient if the second level regions are maximally exhaustive subsets of the outer regions.

Proof. Because the maximum membership subset of a triplet factor is a node pair, the maximum interaction between any two triplet outer regions is a pair-wise interaction. The maximally exhaustive property ensures the number of local pair-wise interactions is at the maximum. Thus, interaction strength among all outer regions is sufficient in the sense of the number of interactions. \square

To satisfy Theorem 2 we first need to generate local pair-wise interactions, which is achieved by defining two types of outer region members for our BFG models: primary triplets and interaction triplets.

Definition 4. A *Primary Triplet* $\mathcal{F} = (\mathcal{V}_{X_i}, \phi)$ is a triplet with nodes set $\mathcal{V}_{X_i} = \{X_i, X_j, X_p\}$ in the moral graph $M[G']$ and a factor ϕ defined by the conversion from the CPD $P(X_i|X_j, X_p)$ in the BFG, G' , as a child variable X_i and its two parents X_j and X_p .

Definition 5. An *Interaction Triplet* $\mathcal{U} = (\mathcal{V}, \phi)$ is a triplet with factor ϕ defined as a uniformly distributed factor, and triplet nodes $\mathcal{V} \in M[G']$ where \mathcal{V} is not equal to any primary triplet's nodes set.

We also need to define:

Definition 6. A *Maximum Membership Subset* of a primary triplet $\mathcal{F} = (\mathcal{V}_{X_i}, \phi)$, $\Omega_{X_i}(X_i \text{ is child node of } X_j \text{ and } X_p)$, is the set of combinations of all node pairs in $\mathcal{V}_{X_i}: \{X_i, X_j\}, \{X_i, X_p\}$ and $\{X_j, X_p\}$.

The shared nodes between two primary triplets will mostly contain a single node only and this obviously fails to satisfy Theorem 2. Instead, by adding interaction triplets we can use node pairs that belong to different primary triplets, which creates a maximum membership subset via which two or more primary triplets can interact. This method of adding new regions to create interactions is also evident in [27].




Fig. 4. (a) κ_4 BFG; (b) factor graph of (a), with factor $f_1 = \phi_{X_1}$ and $f_2 = \phi_{X_1 X_2}$ multiplied into f_3

Figure 4 shows an example where pair-wise interaction is identified using only primary triplets factors, but does not meet the maximally exhaustive property. The factor graph shows a cycle propagation on X_1 and X_2 . This means X_1 and X_2 are strong pair-wise correlated and can be identified as a pair-wise interactions connecting $\{X_1 X_2 E_1\}$

9. A Markov Network is a set of random variables having a Markov property described by an undirected graph.

10. Because a factor in undirected models is defined on all variables it contains.

and $\{X_1X_2X_3\}$ primary triplet regions. However, it is not explicitly known to what extent the node pair $\{X_2, E_1\}$ (a maximum subset of $\{X_1X_2E_1\}$) is also pair-wise correlated, since $\{X_2, E_1\}$ is not cyclicly propagated between factors and is not shared by any two primary triplets. This problem of lack of interaction can be fixed by adding an interaction triplet region $\{X_2X_3E_1\}$. Similarly, other pair-wise correlations will also be incorporated by adding other interaction triplets.

Hence, second level regions in our region graph will be an exhaustive set of all possible pair-wise interactions among all triplet factors.

Although Theorem 2 provides a sufficient condition for the number of interactions, there are also other necessary conditions required for the accuracy of a region graph, as Theorem 3 shows.

Theorem 3. A region graph that does not satisfy both the perfect correlation property and maxent-entropy normal property will not be computationally accurate [9].

Theorem 3 has been informally proved in [9] and is usually used as the necessary conditions to guide region choice.

Next, our ORI algorithm will first identify all outer regions that satisfies Theorem 2 and then identify some redundancy regions to reject, in order to satisfy Theorem 3. As all primary triplets factors are CPD conversions from a BFG, they are already identified in the moral graph. All we need to do next is to identify the interaction triplets and this is achieved by using a coupled Markov Blanket¹¹

Definition 7. A coupled Markov Blanket for nodes (X_i, X_j) is the set of nodes $\partial(X_i, X_j)$ composed of X_i and X_j 's Markov blanket excluding nodes (X_i, X_j) . Therefore $\partial(X_i, X_j) = \partial X_i \cup \partial X_j$ and $\partial(X_i, X_j) \cap (X_i, X_j) = \emptyset$.

The coupled Markov Blanket limited the number of candidate nodes that will be used to generate candidate interaction triplets for a node pair $\{X_i, X_j\}$. The Markov property encoded by a coupled Markov Blanket ensures that our candidate interaction triplet is optimally localized (considered CI information) to capture the local pair-wise correlations.

In Algorithm 1 (ORI) each node pair $\{X_a, X_b\} \in \Omega_{X_i}$ has a coupled Markov Blanket $\partial(X_a, X_b)$. Interaction triplets $\{X_a, X_b, X_c\}$ are then identified for each node pair $\{X_a, X_b\}$ and each node X_c in $\partial(X_a, X_b)$, $X_c \in \partial(X_a, X_b)$ ($a \neq b \neq c$).

ORI ensures all local pair-wise correlations are incorporated by introducing interaction triplets to exhaust all possible local pair-wise interactions (including a node pair that is connected as a moral edge in $M[G']$). So the conditional dependence information, imposed during conversion of the BN to the corresponding MN, is incorporated. Each interaction triplet is initialized with uniformly distributed factor.

11. Markov Blanket for a node X is a set of BN nodes that is composed of its parents, its children and its children's other parents, to guarantee conditional independence between nodes inside and outside of the set [3].

Algorithm 1: ORI algorithm

Input: p' factorized by G' , $p'(x) = \prod_{i \in V} P(X_i | pa\{X_i\})$
 $= \prod_{i \in V} \phi_{\{X_i\} \cup pa\{X_i\}}(X_i, pa\{X_i\})$

Output: Interaction triplets \mathcal{U}

Initial: Interaction triplet $\mathcal{U} \leftarrow \emptyset$;
all primary triplets $\mathcal{F}_i = (\mathcal{V}_{X_i}, \phi_i)$ by;
 $\mathcal{V}_{X_i} \leftarrow \{X_i\} \cup pa\{X_i\}$;
 $\phi_i \leftarrow \phi_{\{X_i\} \cup pa\{X_i\}}(X_i, pa\{X_i\})$;

for each \mathcal{F}_i , the maximum subsets $\Omega_{X_i} \subset \mathcal{V}_{X_i}$ **do**
 for each node pair $(X_a, X_b) \in \Omega_{X_i}$ **do**
 for each node $X_c \in \partial(X_a, X_b)$ **do**
 $\mathcal{U} \leftarrow \mathcal{U} \cup \{(X_a, X_b) \cup X_c\}$ ($a \neq b \neq c$);

for each interaction triplet $\mathcal{U}_i \in \mathcal{U}$ **do**
 if \mathcal{U}_i contains a node pair that is not directly connected in $M[G']$ **then**
 Reject \mathcal{U}_i ;
 if \mathcal{U}_i contains a node pair that is a moral edge in $M[G']$ and \mathcal{U}_i contains root node of G' **then**
 Reject \mathcal{U}_i ;

return \mathcal{U} ;




Fig. 5. (a) κ_5 BFG with a moral edge shown as dashed line; (b) κ_5 BFG with directions removed and all moral edges shown as dashed lines; (c) moral graph of a κ_5 BFG

Figure 5 (a) shows a primary triplet $\{X_1X_2X_3\}$ and an interaction triplet $\{X_1X_3E_1\}$ interacted via an edge (X_1, X_3) .

In Figure 5 (b), for the primary triplet $\{X_1X_2X_3\}$, all its maximum subsets (or corresponding edges), (X_1, X_2) , (X_1, X_3) and (X_2, X_3) , which interact via primary triplet $X_1X_2X_3$ and its interaction triplets. For example, to identify interaction triplets for edge (X_1, X_3) (shown as a bold solid line), we first identify the coupled Markov Blanket $\partial(X_1, X_3) = \{X_2, E_1, E_2, X_4, E_3\}$ (shown with nodes shadowed). From this the candidate interaction triplets are easily identified as: $\{X_1X_3E_1\}$, $\{X_1X_3E_2\}$, $\{X_1X_3X_4\}$, $\{X_1X_3E_3\}$ and $\{X_1X_3X_2\}$. Notice that $\{X_1X_3X_2\}$ is then excluded as it is a primary triplet. Repeated interaction triplets resulting from the selection of other edges' candidate interaction triplets will be removed. All primary triplets can be explicitly identified in the moral graph $M[G']$, and all maximum subsets, as node pairs, are identified as the edges in $M[G']$ as shown in Figure 5 (c). However, by selecting interaction triplets within a coupled Markov blanket we also introduced redundant triplets, which we will subsequently identify and remove.

4.2 Redundancy of interaction triplets

In Theorem 3 we introduced the perfect correlation property which requires that the sum over all counting numbers should be one. We identify two kinds of redundant interaction triplet regions that need to be rejected during outer region identification to satisfy this perfect correlation property:

Type 1, An interaction triplet that contains a node pair (edge) that does not exist in $M[G']$;

Type 2, An interaction triplet that contains a root node of G' and a moral edge in $M[G']$.

Both types can be removed from the collection of all interaction triplets (as will be described in Proposition 3 below).

Because an interaction triplet $\{X_i, X_j, X_k\}$ is defined to connect two or more primary triplets via their maximum subset (node pairs) the redundancy of an interaction triplet can be determined by testing if the entropy of any variable in $\{X_i, X_j, X_k\}$ changes after introducing interaction triplet $\{X_i, X_j, X_k\}$, given other conditions are fixed.

To consider the entropy for each variable in $\{X_i, X_j, X_k\}$ we first need to determine the relationship among the three node pairs, in order to analyse the relationship among variables X_i, X_j and X_k .

The relationship among node pairs can be quantified by the mutual information¹² of node pair region beliefs. There are three node pairs in $\{X_i, X_j, X_k\}$ and there is at least one node pair that is equivalent to a corresponding moral edge or that is not directly connected as an edge in $M[G']$. This means such a node pair $\{X_j, X_k\}$ has a pair-wise factor $\phi_{j,k} = 1$. This uniform factor encodes only conditional independence information $j \perp k | pa\{j, k\}$ and hence X_j and X_k can be separately considered via other node pairings. So the mutual information between this node pair region to another node pair region is equal to the entropy over the shared variable. So the entropy over X_j and X_k will be determined by other two node pairs ($\{X_i, X_j\}$, $\{X_i, X_k\}$) after introducing $\{X_i, X_j, X_k\}$, given other conditions are fixed.

Next we only need to consider the mutual information of the other two node pairs ($\{X_i, X_j\}$, $\{X_i, X_k\}$) which have non-uniform factors. Suppose $b_{i,j}$ and $b_{i,k}$ are region beliefs associated with region $\{X_i, X_j\}$ and $\{X_i, X_k\}$, given $j \perp k | pa\{j, k\}$ we have:

$$\begin{aligned} I(b_{i,j}; b_{i,k}) &= H(b_{i,k}) - H(b_{i,k} | b_{i,j}) \\ &= H(b_{i,k}) - H(\sum_i b_{i,k}) \\ &= H(\tilde{b}_i) \end{aligned}$$

where $H(\tilde{b}_i)$ is the entropy of the marginal belief over variable i in these two regions containing the node pairs.

In summary, $H(\tilde{b}_i)$ is mutual information that in turn determines the entropies of the two node pair regions. So we can reject an interaction triplet $\{X_i, X_j, X_k\}$ by determining if $H(\tilde{b}_i)$ is changed by introducing $\{X_i, X_j, X_k\}$ given other conditions are fixed.

12. Region entropy is defined as $H(b_r) \equiv -\sum b_r(x_i) \ln b_r(x_i)$, mutual information for two region beliefs are $I(b_r; b_s) = H(b_s) - H(b_s | b_r)$.

Based on the above derivation over $\{X_i, X_j\}$ and $\{X_i, X_k\}$ we can prove:

Proposition 3. Type 1 and Type 2 triplets are redundant.

Proof. We have an interaction triplet $\{X_i, X_j, X_k\}$ connecting primary triplets $\{X_i, X_j, X_p\}$ and $\{X_i, X_k, X_q\}$ to determine $H(\tilde{b}_i)$ (X_i is short for i in the following), where $j \neq k$.

We first show Type 1 redundancy which occurs only when $p \neq q$:

Given $p \neq q$, $\{i, j\} \cap \{i, k\} = \{i, j, p\} \cap \{i, k, q\} = \{i\}$, the shared subset of two primary triplets is a single variable i . Based on Equation 4 and the definition of entropy, after message calibration results in Equation 7:

$$\begin{aligned} H(\tilde{b}_i) \propto \tilde{b}_i &= \sum_{j,p} \tilde{f}_{i,j,p} \prod m_{i,j} \prod m_{i,p} \\ &= \sum_{k,q} \tilde{f}_{i,k,q} \prod m_{i,k} \prod m_{i,q} \end{aligned} \quad (7)$$

where $\tilde{f}_{i,j,p}, \tilde{f}_{i,k,q}$ are triplet factors associated with the two primary triplets that contain X_i . All pair-wise messages are incoming messages to the two primary triplets, which are sent from the child regions of the two primary triplets.

Messages $m_{i,p}$ and $m_{i,q}$ do not result from the introduction of $\{X_i, X_j, X_k\}$ so they are 'fixed' here. The messages that vary because of the introduction of $\{X_i, X_j, X_k\}$ are $m_{i,j}$ and $m_{i,k}$. Based on Equation 5 and 6, $m_{i,j}$ and $m_{i,k}$ are all messages originated from factor regions to regions for node pairs $\{i, j\}$ and $\{i, k\}$, in the form of pair-wise and singleton messages. Among these factor regions only those contain i and determine local pair-wise correlation of $\{i, j\}$ or $\{i, k\}$, determines $H(\tilde{b}_i)$ given other conditions are fixed. We therefore only need to find the factors that contain i and also contain variables that j or k depends on through i (so it determines local pair-wise correlations of $\{i, j\}$ or $\{i, k\}$). We denote the set of these factors as Φ_i .

If Φ_i is composed by the two primary triplet factors exclusively, we can reject the introduced interaction triplet, since the local pair-wise correlation over $\{i, j\}$ or $\{i, k\}$ are encoded already in the two primary triplets.

To identify Φ_i we use Figure 6 below.




Fig. 6. (a) $\{i, j, k\}$ not redundant; (b) $\{i, j, k\}$ type 1 redundant; (c) $\{i, j, k\}$ not redundant; (d) $\{i, j, k\}$ type 2 redundant

Figure 6 illustrates partial structures of a BFG G' , with the node pair that is either a moral edge or that is not directly connected as an edge in $M[G']$ is marked with dashed line.

In Figure 6 (a), if both j and k are i 's children, Φ_i is composed by the two primary triplet factors $\tilde{f}_{i,j,p}, \tilde{f}_{i,k,q}$, and other factors that contain i and $pa\{i\}$, such as factor $f_{t,i}$ (there exists paths $t \rightarrow i \rightarrow j$ and $t \rightarrow i \rightarrow k$). So in Figure 6 (a) Φ_i are not determined exclusively by the two primary

triplet factors. Therefore regions for node pairs $\{i, j\}$ and $\{i, k\}$ are necessary in the region graph to capture other factor's information beside the two primary triplets; the interaction triplet cannot be rejected.

If k is i 's child, j must be i 's parent (j, k cannot be i 's parent at the same time), as shown in Figure 6 (b). Here all factors that determine Φ_i are factors in the two primary triplets. The interaction triplet is one that contains a node pair $\{j, k\}$ that is not an edge in $M[G']$ and can be rejected (type 1).

Next, we consider Type 2 redundancy, which only occurs when $p = q$:

Given $p = q$, the only circumstance is when $i = 1$ or 2 , $\{i, j\} \cap \{i, k\} \subset \{i, j, p\} \cap \{i, k, q\} = \{X_1, X_2\}$.

If $X_i = X_2$ (Figure 6 (c)) X_2 must have parent X_1 and apart from factors $\hat{f}_{i,j,p}, \hat{f}_{i,k,q}$, there exists a factor $\hat{f}_{X_1X_2}$ that also determines Φ_i and so $\{i, j, k\}$ cannot be rejected.

If $X_i = X_1$ (Figure 6 (d)), X_1 is root node with a singleton factor \hat{f}_{X_1} , which is always associated to X_1 . Thus Φ_i is determined by \hat{f}_{X_1} and the two primary triplet factors, which will not change by introducing $\{i, j, k\}$. The corresponding redundant interaction triplet is then one that contains a root node and a moral edge (type 2). \square

We use the graph in Figure 5 (c) to demonstrate how proposition 3 is applied.

TABLE 1
Redundant interaction triplet example

Candidate	Node pairs	Primary triplets	Redundancy
$\{X_1X_3X_4\}$	$\{X_1X_3\}$ $\{X_3X_4\}$	$\{X_1X_2X_3\}$ $\{X_3X_4E_1\}$	Type 1
$\{X_2X_3E_1\}$	$\{X_2X_3\}$ $\{X_2E_1\}$	$\{X_1X_2X_3\}$ $\{X_1X_2E_1\}$	N/A
$\{X_1X_3E_2\}$	$\{X_1E_2\}$ $\{X_1X_3\}$	$\{X_1X_2E_2\}$ $\{X_1X_2X_3\}$	Type 2

Table 1 lists three candidate interaction triplets at the first level regions for a region graph built for Figure 5 (c). The related interaction regions (node pairs) and primary triplets are also listed for clarification.

For example, in Table 1, for interaction triplet $\{X_1X_3X_4\}$, the shared subset of both node pairs and primary triplets is the singleton $\{X_3\}$, and X_3 does not have any parent from other factors except these two primary triplets. So this interaction triplet is an instance identical to Figure 6 (b).

Likewise, in table 1 interaction triplet $\{X_2X_3E_1\}$ has node pairs $\{X_2, X_3\}$ and $\{X_2, E_1\}$ with shared subset $\{X_2\}$, and the primary triplets including these node pairs share a subset $\{X_1, X_2\}$. This interaction triplet cannot be rejected as it is an instance of Figure 6 (c).

The interaction triplet $\{X_1X_3E_2\}$ can also be rejected as it is an instance of Figure 6 (d).

4.3 Verification of Theorem 3 for TRC region graph

Now all outer regions are determined by primary triplets plus interaction triplets the corresponding region graph can be generated by the CVM algorithm. The resulting region graph for our BFG models contain three levels,

with all first level region counts equal to one (as all factors are included in first level). The resulting region graph is our TRC region graph which now show will satisfy both necessary conditions of Theorem 3.

In general the TRC region graph properties are summarized in Table 2 (a Proof of these results is given in Appendix A).

TABLE 2
Properties for κ_n ($n > 3$) dimensional $\mathcal{G}(G')$

Levels	$v(r)$	$length$	$max(c_r)$	$min(c_r)$
1 st level	3	$(n-2)^2$	1	1
2 nd level	2	$(n-2)^2$	-1	$3-n$
3 rd level	1	$(n-3)$	$n-3$	1

Table 2 illustrates the region size $v(r)$, the number of regions contained in each level, and max and min of counting numbers in each level's regions.

Proof that the TRC region graph satisfies the perfect correlation property [9]:

Proof. Based on Table 2 (and Appendix A), the sum of all first level region counts is $(n-2)^2 \times 1 = n^2 - 4$. The second and third level regions' counts are cancelled by each other, which will leave one region with counting $3 - n$ (there are two regions at the second level with counting $3 - n$ and one is cancelled) and $(n-2)^2 - (n-3) - 1$ regions with counting -1 . So, sum them all to obtain $n^2 - 4 + 3 - n + ((n-2)^2 - (n-3) - 1) \times -1 = 1$. \square

Proof that the TRC region graph satisfies maxent-entropy normal property [9]:

Proof. The Bethe approximation is maxent-normal [9], and so the entropy of the region graph, H_G , can be written as $H_G = \sum_{i=1}^N H(b_i) - \sum_{a=1}^M I(b_a)$ where N is the number of variables in the region graph, X_i , and M is the number of factors, a , (\mathbf{x}_a are the variables defined by the factor a). $H(b_i) \equiv -\sum_{x_i} b_i(x_i) \ln b_i(x_i)$ is the sum of entropies from all variables X_i in the region graph, and $I(b_a) \equiv \sum_{\mathbf{x}_a} b_a(\mathbf{x}_a) \ln b_a(\mathbf{x}_a) - \sum_{i \in N(a)} H(b_i)$ is the mutual information which is the entropy for a region containing factor a , minus the entropies of all variables contained

in factor a . H_G is maximal, equalling $\sum_{i=1}^N H(b_i)$, when all beliefs, $b_i(x_i)$ and $b_a(\mathbf{x}_a)$, are uniform, and under these circumstances the mutual information, $I(b_a)$, equals zero. In our region graph we can always construct H_G in the form of

$H_G = \sum_{i=1}^N H(b_i) - \sum_{a=1}^M I(b_a)$ because the mutual information for each triplet can be constructed by its connected second level regions and the single variables the triplet contains, resulting in minimal I terms and maximal entropy H_G when all beliefs are uniform. The rest of the proof is omitted for brevity because the verification can be done directly on the TRC region graph. \square

An example of TRC region graph for Figure 5 (c) is shown in Figure 7.




Fig. 7. TRC region graph for Figure 5 (c)

4.4 Region Graph Binary Factorization Algorithm

To use message updating equations on the region graph we need to solve the numerical instability problem discussed in section 3.2. Recall that $c_r = 1 - \sum_{r' \in \text{Ancestor}(r)} c_{r'}$, and so a large absolute value of counting number also implies global multiplicity of connected regions. In Table 2, the connections between first and second level regions grow because the $\min(c_r)$ is linearly decreasing (conversely $\max(c_r)$ is linearly increasing), which means the number of multiple connections grow and we are guaranteed to encounter a numerical instability problem from multiple cycles in the region graph. This is also evident in the example we have used in Figure 7. The CCCP algorithm suffers from the same problem. To reduce the absolute value of the counting number and decompose the multiple connections within a region graph we use the following RGBF algorithm.

Definition 8. A Region Graph Binary Factorization (RGBF) algorithm is one that ensures that each region in a region graph, originally with more than two parents, has exactly two parents without changing the validity of a region graph.

The particular RGBF that we propose is described in Algorithm 2. This RGBF algorithm will be used to generate an equivalent region graph \mathcal{G}' from the original region graph \mathcal{G} , with the properties described in the following proposition.

Proposition 4. By applying the RGBF of algorithm 2 we transform a k -level CVM region graph \mathcal{G} with all factors included in the 1st level, into an equivalent k -level region graph \mathcal{G}' , such that each region r in \mathcal{G}' ($r \in R, r \notin R_{1st\text{ level}}$) is connected to two parents. The counting numbers for all regions are 1, -1 and 0.

Proof. Algorithm 2 will produce $p_r - 1$ (p_r the number of parents) copies of region r in \mathcal{G} to \mathcal{G}' when c_r is not 1, -1 or 0 in \mathcal{G} . Each r region in \mathcal{G}' will share one parent with its neighbouring r copy. Equivalence between \mathcal{G} and \mathcal{G}' can be proved by satisfying the consistency and unity conditions¹³ for a region graph.

- 1) Consistency: as the first level is not changed, consistency of all r ($r \notin R_{1st\text{ level}}$) and its copies with their parents in \mathcal{G}' must be maintained. This is satisfied as each r is connected with its neighbouring copy by sharing one parent, so all parents and all regions r are connected and hence consistent.

13. Unity is defined as where the sum of all regions counting numbers associated with each variable should be one.

- 2) Unity: global unity for each variable must be the same in \mathcal{G} and \mathcal{G}' . As \mathcal{G}' does not contain any new regions compared to \mathcal{G} but only copies of regions, r , from \mathcal{G} , the counting number for each variable will only be influenced by the region r and its copies. So the unity condition can be satisfied by integer accumulation of r and its copies' counting numbers in \mathcal{G}' to c_r in \mathcal{G} , $\sum_{i=1}^{p_r-1} c_{r_i} = c_r$ ($r_i \in \mathcal{G}', r \in \mathcal{G}$), which will not change the unity condition for each variable. In this way the cumulative counting number is not unique but can be specified by using 1, -1 and 0 as these work for any integer. \square

Algorithm 2: RGBF algorithm

Input: k -level CVM region graph \mathcal{G} with regions R

Output: k -level region graph \mathcal{G}'

Initialize: $\mathcal{G}' \leftarrow \emptyset$;

$\mathcal{G}' \leftarrow \mathcal{G}' \cup \mathcal{G}_{1st\text{ level}}$;

for $i = 2 : k$ **do**

for each region $r \in R_{i\text{th level}}$ **do**

if $p_r > 2$ **then**

for $z = 1 : p_r - 1$ **do**

$r'_z \leftarrow \text{repeat } r$

 connect r'_z to two $\text{parent}(r)$

$c_{r'_z} \leftarrow \text{cumulative integer total } c_r$

$\mathcal{G}_{i\text{th level}}$ replace r by r'_1, \dots, r'_{p_r-1}

$\mathcal{G}' \leftarrow \mathcal{G}' \cup \mathcal{G}_{i\text{th level}}$

return \mathcal{G}' ;

The benefit of applying the RGBF algorithm on a region graph is that large counting numbers no longer occur and multiple connections are decomposed into local connections. Therefore, the number of cycles in the region graph is reduced to a minimal number and each third level region will be associated with at most one cycle. An example is shown in Figure 8.




Fig. 8. (a) region graph \mathcal{G} (all 1st level regions' counting numbers 1); (b) region graph \mathcal{G}' by RGBF process of (a)

In Figure 8, regions e and h are copied twice and three times respectively. The counting numbers for each 2nd level region becomes -1, and for each 3rd level region becomes 0 or 1. It does not matter if 1 is placed on h_1 or h_2 since it does not change the consistency and unity conditions, but it will influence the convergence speed. The RGBF process does not change the consistency and unity conditions indicated by \mathcal{G} . If we used the GBP algorithm there would be a limited number of messages multiplied into equations 4 and 6 at each updating iteration, and there is no large counting number for calculating b_r .

If we use RGBF algorithm for a multiply connected CVM region graph, the CCCP updating will also be robust.

CCCP update rule: [20]

$$h_r(x_r) = e^{-\frac{c_r}{c_{\max}}\{E_r(x_r)+1\}} \{b_r(x_r)\}^{\frac{c_{\max}-c_r}{c_{\max}}} \quad (8)$$

$$g_r(x_r) = e^{-\gamma_r - \sum_{s \in \text{child}(r)} \lambda_{r \rightarrow s}(x_s) + \sum_{v \in \text{parent}(r)} \lambda_{v \rightarrow r}(x_r)} \quad (9)$$

$$b_r(x_r) = h_r(x_r)g_r(x_r) \quad (10)$$

$$e^{2\lambda_{r \rightarrow u}(x_u; \tau+1)} = e^{2\lambda_{r \rightarrow u}(x_u; \tau)} \frac{\sum_{x \in r \setminus u} b_r}{b_u} \quad (11)$$

where λ and γ are parent-child region consistency and normalization Lagrangian multipliers. c_{\max} is the max value of all regions' counting numbers in a region graph. h_r and g_r are pre-calculated parameters for computing belief term b_r .

In the CCCP algorithm, updating each $\lambda_{r \rightarrow u}$ is a recursive process that involves calculating the beliefs over all u 's parents and children, and its children's parents. In Figure 8 (a) updating $\lambda_{a \rightarrow e}$ involves the belief calculations of seven regions at a time: a, b, c, e, f, g and h . This number grows with the multiple connections for h and the number of cycles associated with h also grows (there are three cycles associated with h in Figure 8 (a)). But after applying RGBF to (a), as shown in (b), to update $\lambda_{a \rightarrow e_1}$ there are now five regions (a, b, e_1, e_2, h_1) and this number does not increase with the number of connections because there are no multiple connections and only one cycle, maximum, for each level three region.

4.5 TRC Algorithm and Its Complexity

Algorithm 3: TRC algorithm

Input: κ_n dimensional BFG $G', \varepsilon = 1.0e - 5$
Output: G' with marginal distributions
Initialize: $M[G'] \leftarrow$ parametrizing BN G' to MN;
 $\mathcal{F} \leftarrow \phi_j(\phi_j \in p'); \mathcal{U} \leftarrow \emptyset;$
for moral edge $\mathcal{M}_t = (X_i, E_t), (\mathcal{M}_t \in M[G'])$ **do**
 $\mathcal{U} \leftarrow \mathcal{U} \cup \{\mathcal{M}_t \cup X_i\} (X_i \in \partial \mathcal{M}_t);$
 Drop redundancy in $\mathcal{U};$
 $\mathcal{G} \leftarrow CVM(\{\mathcal{F} \cup \mathcal{U}\});$
 $\mathcal{G}' \leftarrow \text{Algorithm2}(\mathcal{G});$
parallel for $r (r \in R_{1st\text{level}}, R_{1st\text{level}} \in \mathcal{G}')$ **do**
 if $|b_r^{\text{old}} - b_r| > \varepsilon$ **then**
 $b_r \leftarrow \text{equation10}$
 if $\text{child}(r) \neq \emptyset$ **then**
 $b_{u \in \text{child}(r)} \leftarrow \text{recursion}(\text{equation10})$
 $\lambda_{r \rightarrow u} \leftarrow \text{equation 11}$
return $G';$

For convenience we can use the CCCP update rule in parallel because to update each λ there are only limited regions involved for computation, which is a result of the

RGBF algorithm. For example, we can update $\lambda_{r \rightarrow u}$ simultaneously when updating $\lambda_{i \rightarrow j}$ provided that $i \neq j \neq r$ and $i \neq j \neq u$. Convergence is guaranteed and is determined by the discrepancy between old and current beliefs. The clustering complexity of TRC is the sum of all levels' regions; this is polynomial and proportional to $\sum_{3 \text{ levels}} v(r) \cdot \text{length}$ (as shown in Table 2) in contrast to exponential clustering complexity for exact methods. Computational complexity is proportional to the number of 1st to 2nd level region edges, which is the sum of all 2nd level's degree of freedom, $\sum_{j=1}^{(n-2)^2} (|c_{r_j}| + 1)$ and is polynomial. A proof of these results is given in Appendix A. Efficiency can be further improved by using parallel processing.

5 EXPERIMENTS

This section presents experiments conducted to determine how well TRC performs for general and high tree-width BFG models compared with competing methods. Typically, experiments carried out in the literature have focused on spin glass models [9], [18], [29]. but we did not use these because they are undirected models and we are interested in directed models. Also, with spin class models the counting number is relatively small and the number of multiple connections is low compared to BFGs. This means spin glass models are easier test cases than BFGs with respect to the numerical instability problem and so we concentrate on experiments involving more challenging BFGs.

Section 5.1 presents the results of testing two simple sparse BN models and compares TRC with the exact and MCMC solutions. Note that, in contrast with MCMC, TRC is not problem tailored and does not need any parameter adjustment.

Section 5.2 presents the results of testing a number of high tree-width BFG models with different numbers of discrete states to investigate efficiency, robustness and accuracy and compare it with a JT solution.

Obviously the RGBF algorithm is an independent step that can be used or not. So in Section 5.3 we compare the results by switching RGBF on/off for both the GBP and CCCP algorithms and find that the RGBF process improves the stability of both algorithms.

Generally, the efficiency of TRC depends on which underlying message update rule is used, either GBP or CCCP or other similar rules, whereas CCCP is slower than GBP [20]. In our tests models of less than 12 dimensions compute within one minute. Computation time for 100 dimensions is up to three hours.

The environment for testing was Java JDK 1.8, Intel E2660 @ 2.2 GHz.

5.1 Sparse BN Model Test

We present the well known Asia model for illustration of a sparse BN graph model by using the TRC algorithm. We change the state names to "1" and "2" to accommodate the result shown in WingBugs [30]. We have restricted our analysis throughout this paper to complete BNs (for reasons explained in Section 2), but as explained in Section 2 we can always convert a sparse BN into an associated complete BN

TABLE 3
Mean value of Asia model variables' marginal results, given variable
 $a = 2$ and $d = 2$

Thre.	Method	s	t	l	b	e	x
	Exact	1.626	1.088	1.100	1.811	1.182	1.220
1.0e-5	TRC	1.626	1.088	1.100	1.811	1.183	1.220
1.0e4	MCMC	1.630	1.084	1.107	1.816	1.185	1.222
1.0e5	MCMC	1.626	1.086	1.100	1.815	1.181	1.218

before applying the BFG process. Hence for sparse BNs to use TRC we use the BFG from its associated complete BN. (see Appendix B for the details).

The TRC result of the Asia model (Table 3) is obtained using a convergence threshold of 1.0e-5. The MCMC threshold, sample size, is increased from 1.0e4 to 1.0e5. Note that in this test TRC is more accurate than MCMC, but more importantly is guaranteed to converge.




Fig. 9. a Dynamic Bayesian Network (DBN) containing three time slices

Using the same approach we also tested the sparse BN model shown in Figure 9. This is a Dynamic Bayesian network (DBN), in the form of the switching state space model [31]. DBN is a popular class of model used in time series analysis, including robotics, protein sequencing and many other domains.

We assume all variables are discretized and are binary variables. Although the model is a sparse BN the tree-width grows with the value of m , where X_m is a hidden vector node. So the exact solution may be intractable when m is large. We set observed values for output node Y in this test. In TRC the clustering complexity remains polynomial regardless of the value of m .

TABLE 4
Mean value of Fig 9 marginal results (with all binary variables and states "1" and "2", NPT setting is in Appendix B)

Method	s_1	s_2	s_3	x_{1t_2}	x_{1t_3}	x_{mt_1}	x_{mt_3}
Exact	1.408	1.782	1.113	1.521	1.797	1.756	1.396
TRC	1.408	1.781	1.114	1.521	1.796	1.756	1.396
MCMC	1.407	1.776	1.112	1.517	1.796	1.762	1.401
MCMC	1.407	1.782	1.114	1.521	1.795	1.758	1.396

Comparison of the exact, TRC and MCMC results as shown in Table 4 (variables x_{1t_1} and x_{mt_2} are not listed as all three methods have the same results). MCMC sample sizes are 1.0e4 (the upper row) and 1.0e5 (the lower row). The TRC convergence threshold is 1.0e-5. TRC approximates all variables very well, achieving a maximum relative error of 0.001 to exact values. Notice that the MCMC result is

improved by increasing the number of samples (from 1.0e4 to 1.0e5), and it achieves a maximum relative error of 0.002 with exact values under the threshold of 1.0e5.

5.2 BFG (from complete graph) Models Test




Fig. 10. κ_{20} BFG (binary variables with random NPTs) test of TRC using different convergence thresholds

We first present a κ_{20} BFG model test by using different convergence thresholds. In Figure 10 we change the convergence threshold ε from 1.0e-3 to 1.0e-5. For each threshold the graph shows the accuracy (measured by the KL distance) compared to the exact solution. As the convergence threshold increases from 1.0e-3 to 1.0e-5 the accuracy increases. When $\varepsilon = 1.0e - 5$ the maximum KL distance is below 1.6e-4, indicating effective improvement of accuracy by increasing the convergence threshold.

TABLE 5
Space complexity and KL distance of TRC for high tree-width BFG models (binary variables with random NPTs) test

	Tree-Width			
	19 (κ_{20})	39 (κ_{40})	79 (κ_{80})	99 (κ_{100})
Space complexity				
JT $O(2^n)$	8 Mb	8e3 Gb	9e15 Gb	9e21 Gb
TRC $O(n^2)$.06 Mb	.11 Mb	.47 Mb	.73 Mb
(iterations)	(782)	(1567)	(3561)	(3963)
KL for TRC				
$\max(KL)$	1.53e-4	1.9e-5	3.8e-5	2.8e-5
$\min(KL)$	3.65e-12	2.5e-13	2.1e-9	4.5e-8
$\text{average}(KL)$	1.46e-5	5.2e-6	7.5e-6	2.9e-6

Table 5 is a summary of the test results for κ_{20} , κ_{40} , κ_{80} and κ_{100} respectively (with $\varepsilon = 1.0e - 5$), tree-width parameter for each model is measured by JT solution. $O(2^n)$ and $O(n^2)$ are space complexities. These results show that the clustering complexity is reduced from exponential to polynomial (from gigabytes to less than 1 megabyte). As the dimensions increase the accuracy does not notably decrease; all KL statistics show a robust and accurate performance and we can increase the convergence threshold to obtain higher accuracy. Because exact computation for all variables is not possible with finite memory, we could not compute exact values beyond 25 dimensions. Therefore, we compare the accuracy of the first 20 dimensions produced under TRC, for all sizes of models, with these exact values and infer that if the TRC results are accurate for these first 20 then the other variables, that could not be directly compared, must be accurate too (because the model has converged).

TABLE 6

κ_{20} BFG model tests with random NPT and different number of discrete states, with convergence threshold set by $1.0e-5$ except the last column

	Discrete state				
	$m = 3$	$m = 4$	$m = 5$	$m = 6$	$m = 6$ ($1.0e-6$)
Iterations	860	642	705	588	1048
KL for TRC					
$max(KL)$	1.15e-5	1.42e-5	1.19e-5	3.14e-5	5.87e-6
$min(KL)$	4.71e-8	1.35e-8	3.18e-8	1.64e-7	2.76e-9
$average(KL)$	2.99e-6	3.9e-6	3.31e-6	7.25e-6	1.12e-6

We compared the TRC result with the exact value for 1 to 10 dimensions for each test model and discovered that as the number of discrete states m , for each variable, increases, the number of iterations required to converge increases in general but also depends on BN parameterizations. This is shown in Table 6, where when convergence threshold is $1.0e-5$, the KL statistics degrade slightly with the increase in the number of states, but when the convergence threshold is set to $1.0e-6$, the KL is reduced. So, as the number of discrete states increases we can set a higher converge threshold to guarantee accuracy.

5.3 RGBF Test




Fig. 11. (a) GBP with/without RGBF on $\varepsilon = 1.0e-5$; (b) CCCP with/without RGBF on $\varepsilon = 1.0e-5$

To investigate the effectiveness of the RGBF algorithm, Figure 11 shows the results of using the GBP and CCCP inference algorithms (showing the problem occurring after $n = 10$) with and without RGBF. Without RGBF both algorithms demonstrate significant inaccuracy (for GBP these are evident in the lowest seven dimensions, and for CCCP these are evident for the lowest four dimensions). In contrast, both algorithms are accurate in all cases with RGBF. These tests also show that the low dimensional variables of a BFG model are more likely to experience numerical problems than high dimensional variables, since low dimensional variables are connected to more children than high dimensional variables.

6 CONCLUSION AND FUTURE WORK

We have presented a general purpose approximate Bayesian Network inference algorithm-Triplet Region Construction (TRC) that overcomes the computational complexity barrier of exact algorithms (such as Junction Tree). Specifically, whereas exact algorithms are worst case exponential for BNs

with large numbers of densely connected variables (since clustering complexity can grow exponentially with the number of nodes) the TRC algorithm reduces the clustering complexity from worst case exponential to polynomial for factorized models. Likewise, the computational complexity is polynomial and can be further speeded up by parallel processing.

The TRC algorithm is based on a binary factorization algorithm and composed of three sub algorithms (ORI, RGBF and CCCP) that provide systematic improvements to previous methods of region based approximate belief propagation (namely relating to region choice, convergence and accuracy). The ORI and RGBF algorithms are independent for high dimensional model problems and can be applied separately to many other types of models.

Experiments carried out by using synthetic data without extreme conditional probability (probability near zero [32]), have shown that TRC is accurate and robust, and so can be used as an alternative to the JT algorithm for handling high dimensional (also high tree-width) problems. Unlike MCMC solutions, TRC is guaranteed to converge and does not require special considerations of parameter adjustment for any discrete models.

Future extensions of this work will focus on using TRC for high tree-width model parameter learning and sensitivity analysis. We will also combine TRC with discretization [12], [14] or sampling for continuous variables so that TRC can take place for all kinds of distributions ¹⁴.

APPENDIX A PROOFS

Proof. of the results in Table 2 of the paper.

- 1) Let n be the number of original nodes in a BFG, G' , so the number of intermediate nodes in G' is: $1 + 2 + \dots + n - 3 = (n - 2)(n - 3)/2$, $n > 3$.
- 2) From the parent to child relationships in G' , the number of primary triplets is determined by the sum of the number of original variables and intermediate variables minus 2, as there are two factors absorbed in triplets. So we have $n - 2 + (n - 2)(n - 3)/2$ primary triplets.
- 3) The number of interaction triplets is the number of moral edges and it is also the number of intermediate nodes, so we have $(n - 2)(n - 3)/2$ interaction triplets.
- 4) The number of first level triplets is then: $L_1 = n - 2 + (n - 2)(n - 3)/2 + (n - 2)(n - 3)/2 = (n - 2)^2$.
- 5) The number of second level intersections is determined by the number of first level triplets and is $(n - 2)^2$.
- 6) There are $n - 3$ intersections with the form $X_i X_j$ which has counting number -1 to $3 - n$, so $\min(c_r) = 3 - n$ at the second level. All other intersections with the form $X_i E_t$ have counting number -1.
- 7) The third level regions are all single variable regions and are original variables X_i , with the counting number 1 to $n - 3$ sequentially, so $\max(c_r) = n - 3$. \square

14. There are preliminary works available in [33].

Proof. of TRC complexity

For all BFG BNs we considered, space complexity is proportional to the sum of all the levels' regions in a TRC region graph. When RGBF is not used, the space complexity is proportional to $\sum_{3levels} v(r) \cdot length$ (as shown in Table 2). When RGBF is used, RGBF results in a linear expansion of second and third level regions so the overall space complexity remains polynomial and proportional to $\sum_{3levels} v(r) \cdot length$.

Efficiency complexity is proportional to the number of first to second level region edges, which is the sum of all second

level's degree of freedoms $\sum_{j=1}^{(n-2)^2} (|c_r| + 1)$ and it is polynomial. RGBF also results in a linear expansion of second level regions, and each second level region has at most four child regions by RGBF, so the second to third level region edges are linear expansions with the number of second level regions. Hence it is linearly proportional to the number of second level regions. Complexity is then determined by the number of first to second level region edges and is still proportional to $\sum_{j=1}^{(n-2)^2} (|c_r| + 1)$. \square

APPENDIX B EXAMPLES

We can add edges to convert a sparse BN to a complete BN before applying the BF process, or alternatively, if the original sparse BN is already a subset of a BFG (i.e. each node has less than two parent nodes) we can also convert this subset of BFG G directly to a BFG G' by using these steps.

- 1) The number of original variables of a corresponding BFG G' equals the number of all variables in G . So if G contains n variables the corresponding G' is a κ_n BFG, in which there are n original variables $\{X_1, \dots, X_n\}$. So there exists a unique parent-child path in G' that contains all original variables of G , which has a parent-child ordering $\pi_{G'} =: \{X_1 \rightarrow \dots \rightarrow X_n\}$.
- 2) Define a parent-child ordering π_G for the n original variables in G , in which for any X_j that is successor to X_i ($i, j \in n$), $X_j \notin pa\{X_i\}$.
- 3) Ensure $\pi_{G'} = \pi_G$ in G' .
- 4) Define NPT for each node X_i ($X_i \in \pi_{G'}$) by reusing the NPT from G , whilst maintaining the CI information encoded in G .

To reuse the NPT for X_i ($X_i \in \pi_{G'}$) from G the parent nodes $pa\{X_i\}$ in G' must be the same as those in G . This is achieved by replicating the original variables through intermediate variables in G' if $pa\{X_i\}$ between G and G' differ.

For example, the Asia model, as shown in Figure 12 (a), is already a subset of a BFG.

For example during the conversion from Figure 12 (a) to (b), the NPT, $P(e|t, l)$ in (a) is reused in (b) as $pa\{e\}$ are the same between (a) and (b). CI information in (a) is also maintained in (b), such as where NPT $P(l|s, E_1)$ is set as $P(l|s, E_1) = P(l|s)$. Here l is independent with E_1 in (b)




Fig. 12. (a) Asia model G ; (b) correct $\kappa_8 G'$; (c) correct $\kappa_8 G'$




Fig. 13. κ_{80} BFG model with node X_1 and X_2 presented

which is identical to (a).

Because $\pi_{G'}$ is not unique we can also obtain alternatives such as that shown in (c). The difference between Figure 12 (b) and (c) for $\pi_{G'}$ is that in (b) $a \rightarrow t \rightarrow s$ and in (c) $a \rightarrow s \rightarrow t$. Although the two alternatives differ both of them contain identical CI information as encoded in the original BN, such as in (b) $P(s|t, a) = P(s)$, and in (c) $P(t|s, a) = P(t|a)$. The marginals that result from the two alternatives are the same.

As there are intermediate variables that are used to replicate the original variables, when setting evidence on the original variable in G' the evidence need also to be set on these replicating variables (marked with the same name as the original variables in (b) (c)).

NPT setting for Table 4

p.s1= c(0.8,0.2),
p.s2|s1= c(0.1,0.9,0.2,0.8),
p.s3|s2= c(0.9,0.1,0.7,0.3),
p.x1t1= c(0.6,0.4),
p.x1t2|x1t1= c(0.7,0.3,0.6,0.4),
p.x1t3|x1t2= c(0.1,0.9,0.4,0.6),
p.xmt1= c(0.3,0.7),
p.xmt2|xmt1= c(0.2,0.8,0.3,0.7),
p.xmt3|xmt2= c(0.4,0.6,0.7,0.3),
p.yt1|s1,x1t1,xmt1 =
c(0.1,0.9,0.2,0.8,0.3,0.7,0.4,0.6,0.5,0.5,0.6,0.4,0.7,0.3,0.8,0.2),
p.yt2|s2,x1t2,xmt2 =
c(0.3,0.7,0.4,0.6,0.5,0.5,0.6,0.4,0.7,0.3,0.8,0.2,0.9,0.1,0.1,0.9),
p.yt3|s3,x1t3,xmt3 =
c(0.6,0.4,0.7,0.3,0.8,0.2,0.9,0.1,0.1,0.9,0.2,0.8,0.3,0.7,0.4,0.6)

Figure 13 illustrates the κ_{80} BFG model we presented in the paper (there are 3083 variables in this BFG). Obviously there are large connections from one node to others, such

as node X_1 and X_2 , which will result in the same node appearing in many different regions in a region graph and introducing numerical unstable problems. So if the first 20 dimensions are accurately approximated, higher dimension variables must be also accurate by the approximation. If otherwise the inaccuracy will appear in low dimensions (i.e. the first 20 dimensions) in priority.

ACKNOWLEDGMENTS

This work is supported by European Research Council Advanced Grant. The full ERC code is ERC-2013-AdG339182-BAYES-KNOWLEDGE.

REFERENCES

- [1] D. Barber, *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.
- [2] F. V. Jensen and T. D. Nielsen, *Bayesian Networks and Decision Graphs*. Springer Publishing Company, Incorporated, 2nd ed., 2007.
- [3] D. Koller and N. Friedman, *Probabilistic Graphical Models - Principles and Techniques*. MIT Press, 2009.
- [4] M. Neil, X. Chen, and N. E. Fenton, "Optimizing the calculation of conditional probability tables in hybrid bayesian networks using binary factorization," *IEEE Trans. Knowl. Data Eng.*, vol. 24, no. 7, pp. 1306–1312, 2012.
- [5] F. R. Bach and M. I. Jordan, "Thin junction trees," in *Advances in Neural Information Processing Systems 14*, pp. 569–576, MIT Press, 2001.
- [6] G. Elidan and S. Gould, "Learning bounded treewidth bayesian networks," in *Advances in Neural Information Processing Systems 21* (D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, eds.), pp. 417–424, Curran Associates, Inc., 2009.
- [7] A. Darwiche, *Modeling and Reasoning with Bayesian Networks*. New York, NY, USA: Cambridge University Press, 1st ed., 2009.
- [8] W. K. Hastings, "Monte carlo sampling methods using markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.
- [9] J. S. Yedidia, W. T. Freeman, and Y. Weiss, "Constructing free-energy approximations and generalized belief propagation algorithms," *IEEE Transactions on Information Theory*, vol. 51, no. 7, pp. 2282–2312, 2005.
- [10] P. Lin, M. Neil, and N. Fenton, "Risk aggregation in the presence of discrete causally connected random variables," *Annals of Actuarial Science*, vol. 8, pp. 298–319, 9 2014.
- [11] A. Gelfand and M. Welling, "Generalized belief propagation on tree robust structured region graphs," in *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence, Catalina Island, CA, USA, August 14-18, 2012*, pp. 296–305, 2012.
- [12] M. Neil, M. Taylor, and D. Marquez, "Inference in hybrid bayesian networks using dynamic discretization," *Statistics and Computing*, vol. 17, no. 3, pp. 219–233, 2007.
- [13] D. Koller, U. Lerner, and D. Anguelov, "A general algorithm for approximate inference and its application to hybrid bayes nets," in *Proc. Fourteenth Annual Conference on Uncertainty in AI (UAI)*, pp. 324–333, 1999.
- [14] A. V. Kozlov and D. Koller, "Nonuniform dynamic discretization in hybrid networks," in *UAI '97: Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence, Brown University, Providence, Rhode Island, USA, August 1-3, 1997*, pp. 314–325, 1997.
- [15] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- [16] T. P. Minka, *A Family of Algorithms for Approximate Bayesian Inference*. PhD thesis, Cambridge, MA, USA, 2001. AAI0803033.
- [17] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky, "Tree-based reparameterization framework for analysis of sum-product and related algorithms," *IEEE Transactions on Information Theory*, vol. 49, no. 5, pp. 1120–1146, 2003.
- [18] V. Kolmogorov, "Convergent tree-reweighted message passing for energy minimization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, pp. 1568–1583, Oct. 2006.
- [19] R. Matescu, K. Kask, V. Gogate, and R. Dechter, "Join-graph propagation algorithms," *J. Artif. Intell. Res. (JAIR)*, vol. 37, pp. 279–328, 2010.
- [20] A. L. Yuille, "Cccp algorithms to minimize the bethe and kikuchi free energies: Convergent alternatives to belief propagation," *Neural Computation*, vol. 14, p. 2002, 2002.
- [21] K. P. Murphy, Y. Weiss, and M. I. Jordan, "Loopy belief propagation for approximate inference: An empirical study," in *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, UAI'99*, (San Francisco, CA, USA), pp. 467–475, Morgan Kaufmann Publishers Inc., 1999.
- [22] A. Braunstein, M. Mézard, and R. Zecchina, "Survey propagation: An algorithm for satisfiability," *Random Struct. Algorithms*, vol. 27, pp. 201–226, Sept. 2005.
- [23] A. Jaimovich, O. Meshi, I. McGraw, and G. Elidan, "Fastinf: An efficient approximate inference library," *J. Mach. Learn. Res.*, vol. 11, pp. 1733–1736, Aug. 2010.
- [24] A. L. Yuille and A. Rangarajan, "The concave-convex procedure," *Neural Comput.*, vol. 15, pp. 915–936, Apr. 2003.
- [25] E. B. Sudderth, A. T. Ihler, M. Isard, W. T. Freeman, and A. S. Willsky, "Nonparametric belief propagation," *Commun. ACM*, vol. 53, pp. 95–103, Oct. 2010.
- [26] A. Ihler, A. Frank, and P. Smyth, "Particle-based variational inference for continuous systems," in *Advances in Neural Information Processing Systems 22* (Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, eds.), pp. 826–834, 2009.
- [27] M. Welling, "On the choice of regions for generalized belief propagation," in *UAI '04, Proceedings of the 20th Conference in Uncertainty in Artificial Intelligence, Banff, Canada, July 7-11, 2004*, pp. 585–592, 2004.
- [28] M. Welling, T. P. Minka, and Y. W. Teh, "Structured region graphs: Morphing EP into GBP," in *UAI '05, Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence, Edinburgh, Scotland, July 26-29, 2005*, pp. 609–614, 2005.
- [29] J. Sun, N. N. Zheng, and H. Y. Shum, "Stereo matching using belief propagation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, pp. 787–800, July 2003.
- [30] D. J. Lunn, A. Thomas, N. Best, and D. Spiegelhalter, "Winbugs; a bayesian modelling framework: Concepts, structure, and extensibility," *Statistics and Computing*, vol. 10, pp. 325–337, Oct. 2000.
- [31] Z. Ghahramani, "Hidden markov models," ch. An Introduction to Hidden Markov Models and Bayesian Networks, pp. 9–42, River Edge, NJ, USA: World Scientific Publishing Co., Inc., 2002.
- [32] P. Dagum and M. Luby, "An optimal approximation algorithm for bayesian inference," *Artificial Intelligence*, vol. 93, no. 12, pp. 1 – 27, 1997.
- [33] P. Lin, *Performing Bayesian Risk Aggregation using Discrete Approximation Algorithms with Graph Factorization*. PhD thesis, 2015.

Peng Lin is PostDoc at School of EECS, Queen Mary U. of London.

Martin Neil is Professor of statistics at School of EECS, Queen Mary U. of London.

Norman Fenton is Professor of Computer Science at School of EECS, Queen Mary U. of London.